The University of Georgia

Franklin College of Arts and Sciences
Department of Psychology

December 22, 1992

AD-A259 860

Susan E. Chipman
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Dear Susan:

Part I of this quarterly report is a description and explanation of two candidate measures of the quality of a person's prior knowledge. It follows up on my brief coverage of these measures in my first quarterly report.

Part II of this present report describes progress on other aspects of the project.

Part I: An Explanation of how Settling Time and Harmony can be used as Measures of Prior Knowledge

After anyone takes a course or reads a text, they end up with a mental representation of it. Our cognitive structure test, which is the focus of the ONR contract, is intended to give us a map of the mental representation of each person who takes the test (or at least it gives us the best available map of it).

For an initial idea of how the cognitive structure test does this, you should know that the cognitive structure test simply asks the learners to tell us the distances between the important ideas in their mental representations. From those distances, we then construct a map depicting the mental representation. The process is identical to what we could do to get a map of the mental representation of each person for some geographical space, say, the United States. We could simply ask them the distances between the important cities in the United States, and then construct from those distances a map depicting their mental representation of the United States.

We did a cognitive structure test for a 1000-word text on the Vietnam War by choosing the 12 most important terms, giving people all possible pairs of the terms, and asking them to tell us the distances between them. The data I'll report below is from that test.

Let's try the same thing for a simple example using cities. Suppose I ask you to tell me the distance between all possible pairs of 3 cities:

--New York and Washington, and you say 500 miles.
--New York and Los Angeles, and you say 3000 miles.
--Washington and Los Angeles and you say 3000 miles.

92

The matrix that was produced is like this:

	New York	Washington	Los Angeles
New York	0	500	3000
Washington	500	0	3000
Los Angeles	3000	3000	0

and the map is like this:
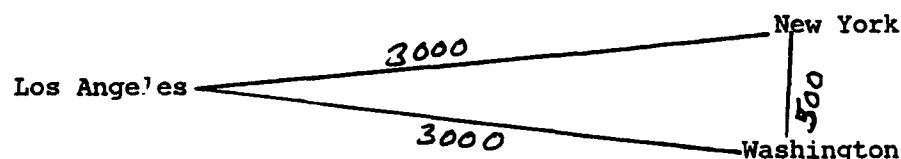
We did this sort of thing for all possible pairs of 12 terms from the Vietnam War text with 7 experts (an ambassador, two military historians, the commander of the local AFROTC, etc.) and 424 Air Force recruits.

The experts, of course, knew a lot about the Vietnam War (that was why we picked them), while most recruits did not. So we presume that the experts gave us a better representation of the text than the novices, just as a U.S. native would give us a better representation of the U.S. map than someone from another country, because the U.S. native has more subject-matter knowledge about U.S. geography than the non-native.

Spreading activation is one way that people think about associated ideas, according to modern cognitive science. If two ideas are strongly associated, one will make you think of the other, i.e., one will excite the other, and activation will spread strongly between them. That is, if you think of one of the ideas, that will probably make you think of the other. But if the two ideas are inhibitorially related, then when you think of one, your likelihood of thinking of the other will decrease. For example, if we think about "Madonna", we become more likely to think about "Sex" because they are excitatorially related. But if we think about "George Bush", we are unlikely to think about Madonna, because they are unrelated or possibily inhibitorially related.

The great advantage of the spreading activation process is that it has been implemented on computers. All we need to know to spread activation is how closely the various ideas are associated, and then we can "think" about them by spreading activation among them. But the thinking is not inside the person's head, where we cannot get at it; instead, it is out in the open in our computer where we can examine it as closely as we want.

I said that all we need to know to implement spreading activation "is how closely the various ideas are associated." But we already know how closely the various ideas are associated, because the cognitive structure test specifically asks how close all the pairs of ideas are. So you can see that we can apply spreading activation to the association matrix that we got from the cognitive structure test.

A-11

So, the main points up to here are:

1. The person creates a mental representation from life experiences, taking a course, or reading a text.

2. We copy that mental representation from the person's mind, using the cognitive structure test.

3. We then study our copy of the person's mental representation, to find out its characteristics.

4. The particular topic we are interested in for the ONR contract is how the person's subject-matter knowledge can be measured by looking at characteristics of their mental representation. We are interested in two aspects of subject-matter knowledge: prior knowledge, and knowledge gained from some sort of learning experience, such as taking a course, or reading a text.

5. To anticipate, the first page of my first quarterly report proposed that two characteristics of the person's mental representation that we can use to measure his subject-matter knowledge are (a) harmony and (b) settling time.

What is Settling Time?

When you think about a set of associated ideas, it takes you a certain amount of time to finish thinking about them. The way you know when you are finished thinking about them is when your thoughts have stopped changing. (If your thoughts have stopped changing, you are not thinking any new thoughts: you have stopped further thinking; the progress of your thinking has stopped.) In the same way, the computer programs that spread activation will "think" about the associated ideas until they are finished thinking about them, and they know when they are finished thinking about them when their "thoughts" have stopped changing.

This is implemented by spreading the activation (from each idea to each idea it is connected to) for one cycle, then measuring the state of activation of all the ideas in the network, and subtracting that state from the state of the ideas on the previous cycle: if there is a change from the previous cycle to the current one, then the program will keep on thinking for another cycle. On the other hand, if there is no change in the state of the ideas from the previous cycle to the current cycle, then the thoughts have stopped changing, so there is no point in thinking further, so the computer has finished thinking, and it stops. It then reports how many cycles it took to stop.

We then use that count of the number of cycles that the computer takes to finish thinking. In our analysis of previously collected data, experts took about 4 cycles for their networks to settle, and novices took about 9 cycles. These differ significantly.

What happened here is that the average subject-matter expert, chosen because they had subject-matter knowledge, settled faster than the average novice, who had less subject-matter knowledge. We think that the experts settled faster __because__ they had more subject-matter knowledge. If so, this means that we can measure people's prior knowledge in an area by measuring the settling rate of their mental representations in the subject matter area. Since the ONR contract is about good ways to measure people's prior knowledge, this finding is

relevant to the contract. However, it should be noted that settling rate is influenced by many other factors and so may not be a suitable individual differences measure. Harmony is a better candidate, and we now turn to it.

What is Harmony?

Harmony is another measure we get when we spread activation around the person's mental representation. It measures the "goodness" of the learner's mental representation. We have found that this measure of the goodness of the mental representation is associated with high levels of subject-matter knowledge, and there are good theoretical reasons why this should be so, so harmony may be a good measure of subject-matter knowledge.

Harmony is determined by the consistency of the set of distances that the learner produces. If the distances all fit together with each other, than high harmony scores will result. If the distances are inconsistent with each other, low harmony scores will result.

In the cities example, the matrix shown before is consistent; a map of it can be drawn on the page. But if the person said that New York was 500 miles from Washington and New York was 3000 miles from Los Angeles, but he also said that Washington was 500 miles from Los Angeles, there is no way you can draw that on the page; the distances are inconsistent.

To measure harmony in the cities example, we would take the association matrix of those distances and spread activation through it. Then we would measure how well the final set of activations met the constraints imposed by the associations that the person gave us. In this case, they couldn't fit very well, since there is no way that Washington can be drawn 500 miles from New York and Los Angeles, and New York can be 3000 miles from Los Angeles, as you can see from this map:

Los Angeles ———— 500 ————— Washington ———— 500 ——— New York

?3000?
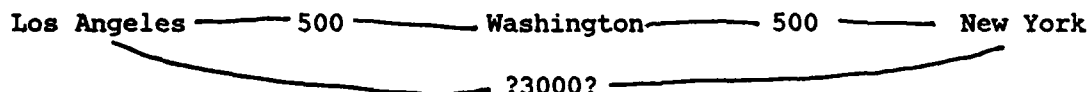
From this point, we have to bid a fond farewell to the cities and distance analogy (since it doesn't help with several important characteristics that spreading activation has) and use a different analogy, which is very similar to the actual procedure used by the spreading activation process.

Suppose there are three nodes, A, B, and C, serving as locations for ideas, and they are associated by links, where the associative links are either excitatory—represented by +1—or inhibitory—represented by -1. If 2 nodes are linked by an excitatory link, then if one is in a certain state, the other wants to be too. But if they are inhibitorially linked, then when one is on, the other will want to be off. We get the links from the subject, as his ratings for the direction and degree of association between each pair of terms.

Some arrangements of such nodes and links have high harmony because they are consistent. For example,

```
        A
       / \
   +1 /   \ +1
     /     \
   B/_____\C
        +1
```
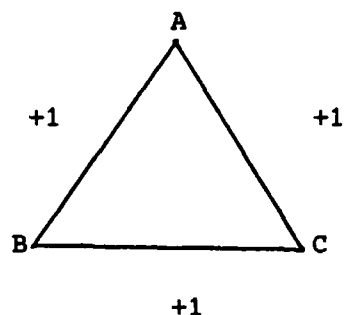
has a harmony score of 1 (harmonies range from -1 to +1). If A is on, B wants to be on because of the +1 on the link, and if B is on, it wants C to be on because of the +1, and if C is on, it wants A to be on because of the +1. So if everything is on, the network will satisfy all the contraints perfectly, and harmony will be at its maximum of 1.

But other arrangements have low harmony because there is no way the network can satisfy them. For example,

```
        A
       / \
   -1 /   \ -1
     /     \
   B/_____\C
        -1
```
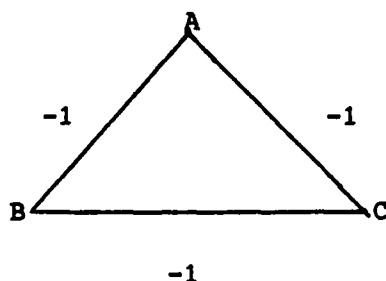
has a harmony of zero for these reasons:

If A is on, it wants B to be off (i.e., in the opposite state) because of the -1 on the associative link. If B is off, it wants C to be on (i.e., in the opposite state) because of the -1 on the associative link. If C is on, it wants A to be off because of the -1 on the associative link. But we started out with A on. So if we turn it off, we would have to change B to on, and C to off, and A to on. But then this would send us on another cycle, endlessly. There is no way the network of ideas can satisfy these constraints.

The two other cases are not so obvious.

```
          A
         / \
     -1 /   \ -1
       /     \
     B/_____\C
          +1
```
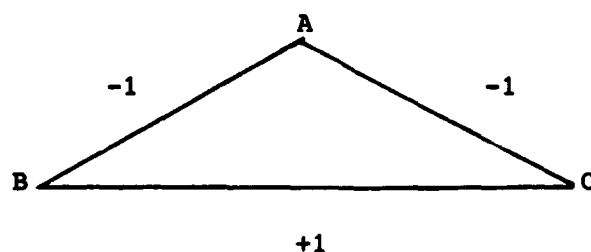
In this case the harmony is 1. This is because if A is on, B wants to be off because of the minus 1. If B is off, then C wants to be off too, because of the plus 1. And if C is off, then A wants to be on, because of the minus 1. So all the constraints reflected in the links are satisfied, and harmony is 1.

But consider this case:
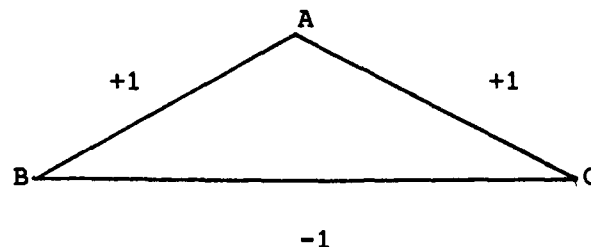
Here the harmony is low, because if A is on, B wants to be on because of the +1; if B is on, C wants to be off, because of the -1; but if C is off, A wants to be off too because of the +1, but A is on. So the constraints are not met.

We do the same sort of thing with the 12 X 12 matrices we got from our subject-matter experts and novices: we arrange them in a network and then spread activation among them. For each person, we get a harmony score. Our results showed that the experts' harmonies ranged from +1 to 0.6. Novices' mean harmony was 0.4. Experts had significantly higher harmony than novices.

What happened here is that the average subject matter expert (chosen because they had high subject-matter knowledge) had higher harmony than the average novice, who had less subject-matter knowledge. We think that the experts had higher harmony because they had more subject-matter knowledge. If so, we can measure people's prior knowledge by measuring the harmony of their mental representations in the subject matter area. Since the ONR contract is directed at good ways of measuring people's prior knowledge, this finding is relevant to fulfilling the contract.

So the upshot is that our harmony and settling time results are consistent with the hypothesis that subject-matter knowledge can be measured by harmony and settling rates.

Effects of Text Quality on Knowledge, Harmony, and Settling Rate

One way to get subject-matter knowledge is to read an understandable text on the subject. Recruits who read an improved text should therefore have more subject-matter knowledge, because the text was improved to be more understandable, and sure enough, we found they have higher harmony and faster settling rates.

This is more evidence consistent with the hypothesis that subject-matter knowledge can be measured by harmony and settling rate.

Individual Differences in Knowledge, Harmony, and Settling Rate

To get at individual differences between recruits in subject-matter knowledge, we divided our 424 recruits into 2 piles: those who had similar cognitive structures to the experts, and those who didn't. Those who had similar structures to the experts presumably had in advance, or had acquired by reading, more subject-matter knowledge by the time they took the cognitive structure test, and sure enough, they had significantly higher harmony and faster settling rate.

This is more evidence consistent with the hypothesis that subject-matter knowledge can be measured by harmony and settling rate.

Summary and Conclusion

Here I summarize the main points, and how they relate to the ONR contract.

1. We measure the person's mental representation by asking them how closely associated are all possible pairs of the important terms in the subject matter area.

2. This gives us an association matrix. In this matrix, we can imagine all the terms listed across the top and along the side, and an entry in each intersection indicating the closeness of association.

Another way to imagine this matrix is as a diagram of an associative network, with each term represented as a node, and each association represented by a link between nodes. Each link can be either excitatory, which leads to a tendency for one node to go on when the other is on, or one node to go off when the other is off; or else inhibitory, leading to a tendency for one node to go off when the other is on.

3. We then use this association matrix to simulate the person thinking about the subject, using the spreading activation program as our working model of the thinking process.

4. After the simulation stops thinking, we measure how long it had to think to reach a stable state, and how well its final state satisfies the constraints in its association matrix. If they match well, it has high harmony.

5. Finally, the critical thing for the ONR contract is that there are relatively few ways for the subject to provide a set of links that satisfy the constraints, especially in a large matrix. One of them--the one we are interested in--is by having a well-developed and therefore consistent view of the subject matter area from which the terms come, and this corresponds to a high level of prior knowledge.

Another way to get a high harmony is by chance, happening to luck into a consistent set of ratings, but this probability is precisely calculable, and decreases precipitously as more terms are used.

Another way is to rate everything as closely related, but this can be spotted easily.

6. Consequently, the harmony measure may be a good measure of the goodness of a person's prior knowledge. It is available from the cognitive structure type

of test, and from no other type of test.

Part II: Progress on other parts of the project

    1. <u>Program Development</u>. We have extended our computer program for interviewing experts to permit construction of a log of their reasons for each response. We have also incorporated a program that can immediately provide a network representation of their responses that they can comment on.

    2. <u>Testing Experts</u>. We have tested several experts and ourselves, and developed several guidelines to tell experts about how to select terms for their structures. These include selecting (a) nouns or noun-phrases, (b) description of states, (c) not selecting verbs, (d) but selecting verb phrases, (e) selecting pairs of evaluative terms like good-bad or success-failure, (f) selecting terms for which experts might have different representations than novices, including (1) jargon terms for which the technical meaning within the area is different than in ordinary language, and (2) technical terms for which no counterpart exists in ordinary language. Guideline f is intended to encourage the selection of terms that will discriminate among experts and novices.

    A preliminary protocol for instructing experts to select such terms has been developed, along these lines: What is the most important thing in this text? What is the next most important thing in this text? Now those two are closely related (similar). Can you pick a thing that is not closely related (similar) to the first one? To the second one? O.K., now can you pick an evaluative term for any of these? And the opposite of that?

    [After printing out all pairs] As you look at these pairs, are there any pairs that you think you would say are closely related (or distantly related) but a novice would not? [Looking now at a list of each term separately] Can you think of a term to go with any of these that an expert would rate in one direction but a novice would rate in the other direction?

    We expect this protocol will be useful for getting experts to choose terms useful for the purposes of the cognitive structure tests.

    3. <u>Theoretical Development</u>. We have developed a model for getting a computer program to make ratings of relatedness, given the propositions of a text or subject-matter area. We are now testing this, using ourselves to provide ratings. This model serves the heuristic function of developing and refining our ideas about how the ratings are generated from the prior knowledge representations.

    I evaluate our progress so far as good. I should note that the first part of this quarterly report was written in response to Wallace Sinaiko's request. If you have any questions or want to talk, please call me at (706_ 542-3094, or at home at (706) 549-4661.

<div style="text-align:center">Sincerely,</div>

<div style="text-align:center">Bruce K. Britton
Professor</div>

BKB/rcm